An Aerial Phytobiopsy System: Design, Evaluation, and Lessons Learned

Daniel Orol, Jnaneshwar Das, Lukas Vacek, Isabella Orr, Mathews Paret, Camillo. J. Taylor, Vijay Kumar

***Abstract*— Early plant disease detection and treatment could dramatically increase crop yield. However, even experts cannot visually distinguish various diseases with certainty, limiting image based diagnosis. In this paper, we present a novel small unmanned aircraft system (sUAS) for phytobiopsy. This platform can remove a leaf section with visual symptoms and transport it to a lab for precise disease analysis. We present the design of a gripper and an arm mechanism, and we discuss best practices for using the system. Results are presented from extensive experimental evaluation of the gripper assembly and 21 indoor manual flight trials, demonstrating efficacy. Our sUAS utilizes its inherent mobility and range to solve a pressing agricultural problem, improving food production and disease detection capabilities. A short video of our system can be found at `https://goo.gl/xomYUO`.**

I. Introduction

Fig. 1: This work addresses phytobiopsy using UAVs that fly close to symptomatic trees and acquire leaves autonomously for ex-situ analysis.

Plant diseases severely impact crop growth and yield, with significant economic impact in the US from citrus greening disease alone. Current detection methods for phytopathology utilize specialized optics [1], as well as manual sample collection, for ex-situ analysis (Fig. 1). However, visual methods are not suitable for precise diagnosis, and manual sample collection does not scale for two reasons. First, many farms are simply too large, resulting in under-sampling. For example, watermelon plots can span up to 100 acres, with scouts having around 10-15 minutes to inspect a farm. Second, once a scout obtains a specimen, she must transport it to a lab equipped with molecular or morphological analysis equipment for precise identification of the disease as well as its stage. Ahlin et. al presented an approach for visual detection and removal of leaves from a plant with an articulated ground robot, closing the loop on visual-servoing for this task [2].

The authors are with GRASP Laboratory, University of Pennsylvania, Philadelphia, PA 19104, USA, {danorol, djnan, vacek, cjtaylor, kumar}@seas.upenn.edu.

*Mathews Paret is at the North Florida Research and Education Center, University of Florida, Quincy, FL 32351, USA, {paret@ufl.edu}.

In this paper, we introduce the first prototype of a novel unmanned aerial vehicle (UAV) platform that enables fast acquisition of leaf samples for ex-situ analysis, mitigating these challenges. Precision agriculture is expected to be a major market for UAS, with expected economic impact over the next few years in billions of dollars [3]. In particular, UAVs meet the requirements of scouting well because they can navigate otherwise inaccessible areas of a farm quickly, while using integrated sensors to maintain process integrity. For example, pecan trees routinely grow 80-100 feet tall, posing hazards for manual inspection, and making them infeasible for surveying with ground robots. Current UAV platforms use onboard imaging systems to survey agricultural farms, enabling 3-D canopy reconstruction, yield estimation, and crop stress detection [4], but few engage in physical interaction. Our system enables access to a currently untapped range of agricultural manipulation and collection tasks.

Aerial manipulation is a rapidly growing field, with systems designed for autonomous grasping [5], [6], perching [7], object manipulation [8], [9], and inspection [10]. Aerial manipulator systems have used various form factors and configurations such as lightweight quadrotors, heavy-lift hexrotors, tilt-rotors, and helicopters [11] [12] [13]. Ore et. al designed an autonomous sUAS for collection of physical samples of water from lakes using a tube attached to the UAV [14]. Lightweight systems are required for these applications due to the inherent payload constraints of aerial systems. The type of aerial contact we deal with generalizes well to challenges such as power cable inspection, construction testing, and any other area requiring mid-flight physical contact. Our work is inspired by the co-robotic design of grippers discussed by Gealy et al [15].

The main contribution of this paper is a lightweight small unmanned aircraft system capable of grasping, severing, storing, and transporting sections of leaves from a plant. Precise positioning of the end-effector with respect to the propellers mitigates UAV downwash. We present the design of the sUAS, the arm system, and the gripper, targeting leaves from citrus, rose, and watermelon plants (Fig. 2). The three plants cover a broad spectrum of leaf size and texture characteristics;

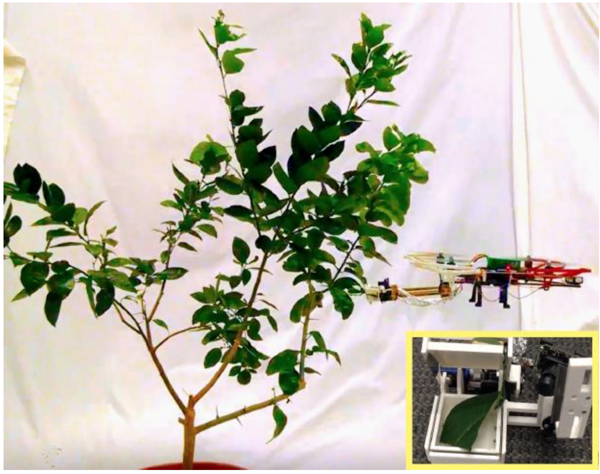

Fig. 2: The system during a flight trial, severing and storing a citrus leaf. Inlay at bottom-right shows the gripper mechanism with the acquired leaf sample. The storage box lid prevents contamination during specimen transport.

therefore, our system can grasp small leaves, sections of large leaves, tough leaves, and soft leaves. Our test results demonstrate a 98.0% success rate for cutting target plant samples, 93.8% for severing a sample leaf during manual flight tests in which the specimen entered the system, and 80.0% for securely storing the specimen properly once severed during manual flight tests. To the best of our knowledge, this work presents the first aerial phytobiopsy system. Our included video shows success and failure cases for manually piloted operation of our system.

The rest of the paper is organized as follows: We describe the system in section II. Extensive evaluation of the gripper assembly, and results of manual flight trials are described in Section III. Finally, we close with a discussion of future work in Section IV.

II. System Description

In developing an aerial robotic platform capable of reliably collecting symptomatic leaves from plants of interest, we identified the following challenges: positioning of an end-effector to mitigate UAV downwash, grasping a leaf, severing it from the plant, and storing it securely for transport (Fig. 3). We iterated through a series of prototypes, exploiting 3-D modeling software and simulations as well as precision 3-D printing and laser cutting (Fig. 4, 5). The cost of our system is under 700 USD, and the total takeoff weight is 1.9 kg. The system is described in detail through the rest of this section.

Explicitly, we set out to design a system capable of aerially collecting and storing symptomatic plant specimens for transport to a lab. This system must be able to stably fly for enough time to collect a specimen and deposit it at the target location. We demonstrate the system's ability to meet these minimum requirements through a series of manual flight tests resulting in collection and storage of a plant specimen collected aerially from a living lime tree. The UAV flew for over 10 minutes without needing a battery recharge.

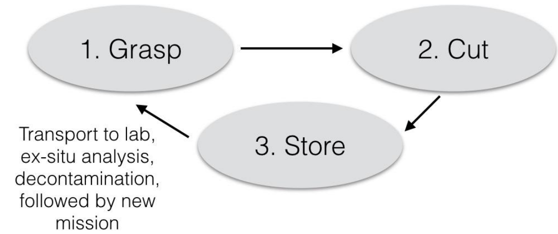

Fig. 3: The three steps in successful leaf acquisition after identification are grasping, cutting, and storing. Following this, the UAV transports the secure leaf to a lab for analysis.

A. Mechanical Design

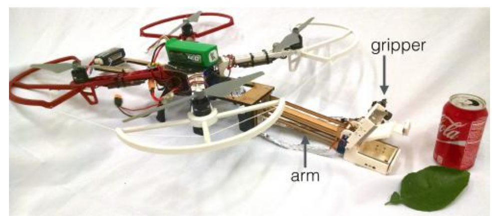

Fig. 4: The phytobiopsy UAV, with the gripper and the arm assemblies highlighted. The soda can is for scale.

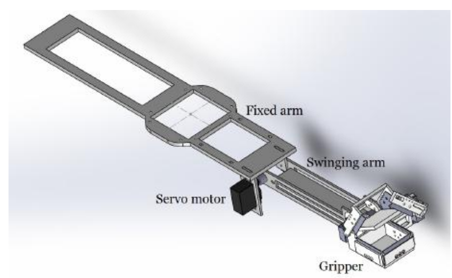

Fig. 5: CAD model of extendable arm, showing the servo motor for extending and retracting the arm.

To keep the system maneuverable and low-cost, we use a DJI F450 airframe with motor-motor diameter of 45 cm, using a Pixhawk flight controller and DJI E310 propulsion system. The gripper design resembles a mouth, with the cutting section forming a 5.3 by 5.3 cm-square receptacle with blades on all three outward facing sides (Fig. 7). This ensures that the leaf or petiole [1] is severed consistently across many leaf angles. Once cut, the sample is stored in a closed box to prevent contamination during transport.

We conducted a set of experiments to qualitatively characterize the downwash in order to determine the ideal gripper placement. The sUAS was affixed using a rigid rod about 2 m above the ground (Fig. 6). The

[1] the plant section connecting the leaf to the stem

Fig. 6: Setup of tests used for characterizing downwash. The sUAS was attached to a ladder, and a plant (bottom left) was manually moved around beneath the robot. We visually observed leaf perturbation from downwash.

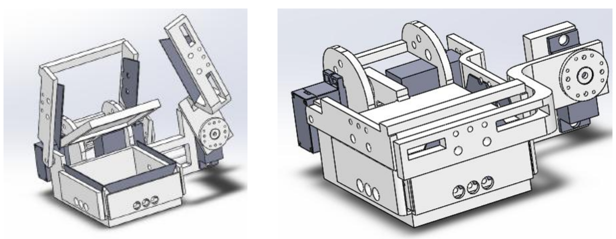

(a) Final iteration in open configuration. (b) Final iteration in closed configuration.

Fig. 7: The end-effector assembly for the phytobiopsy system.

throttle was set to 20%, 30%, 35%, and 40% while a researcher manually moved a branch around beneath the robot. We observed hover, roll, and pitch flight configurations, all of which change the relative thrusts of the motors. The propeller-generated wind visually displaced the leaves beneath the robot, even in the center where there is no propeller. We also tested to see if placing objects in-between the propellers and the plant helped. Three feasible options for end-effector placement were identified – an arm extended out horizontally, a windscreen suspended below the sUAS, or a cable-winch system to extend the gripper beneath the sUAS.

The cable-winch system, although light and versatile, involves considerable mechanical complexity for repeatable motions. More importantly, it presents controls challenges [16]. A well-designed windshield would not affect controller performance; however, this setup requires precise aerodynamic shaping to deflect the air around the plant without altering the thrust strength or motor efficiency.

As opposed to the above two choices, an arm creates an off-center mass but is mechanically simple and controllable [17]. The arm design also provides an extra degree of freedom, rolling 180 degrees with the limits parallel to the UAV's base. This increases the gripper's ability to acquire samples from canopies very close to the ground.

(a) Image of our final iteration holding a leaf. (b) The leaf after cutting operation.

Fig. 8: Demonstration of grasp and cut action on a leaf.

The arm uses a simple, one-joint design (Fig. 5). In flight mode, the gripper will remain beneath the robot so that the system's center of mass is geometrically centered. In sample acquisition mode, the arm will extend, at which point a secondary controller will take hold and stabilize the off-center mass for the short period of time required to acquire a leaf. The gripper is rigidly affixed to the end of the arm. For this paper, we have not yet implemented our sliding mode controller to account for the moving mass. We used the battery that powers the gripper as a counterweight to keep the system stable with a default controller and avoided swinging the arm mid-flight.

B. Grasping

The lid of the storage box, or the *tongue*, is actuated and allows our system to grasp leaves securely. By combining the storage and grasping mechanisms, we maintain mechanical simplicity while enhancing reliability. Our initial prototype used a finger-like end effector (common for grasping), which testing indicated was not as reliable as our current mouth-like design. To keep the grasping and cutting actions separate, the lid actuates separately from the blades, which allows the system to attempt multiple cuts without releasing the sample and provides more consistent results.

In addition to securing the sample, the lid of the storage box creates tension at the cutting interface. This prevents flexible leaves from bending upon contact with the blades, improving the system's performance with regards to both leaves and petioles.

C. Cutting

The system uses stock, 5.3 cm long box-cutter blades. Although these blades, unlike scissors, are tapered on both sides of the tip, we chose them for their low cost, easy maintenance, and successful performance; the custom design of blades is a non-trivial task.

One of the key factors to successful cutting is the spacing between blades. Our prototype leaves 0.127 mm of separation between blades; this distance allows for consistent cuts and lies at the limit of our 3-D printing tolerance. The spacing also led us to implement independent actuation of the front and side blades. If actuated simultaneously, the front blades would have moved along an arc, making it geometrically impossible to maintain consistent spacing. The dual-action mechanism replicates scissors. Our initial tests eliminated parallel actuation of front blades, since this resulted in unreliable leaf tearing. However, when a scissor-like cutting action was used, smooth, reliable cuts were observed (Fig. 8).

To improve reliability and avoid the issue of gear backlash, we chose slightly heavier servos for the blades rather than lighter, weaker servos coupled with gear ratios. Additionally, our analysis showed that we reduced weight by using two servos instead of one heavier servo that sequentially actuated both sets of blades.

The torque for the blades was partially dictated by related research by Mahvash et al. on scissor cutting force [18]. They showed that scissors could cut paper with 2.5–4.5 N of force depending on the size of the scissors. The weakest part of our cutting mechanism can nominally apply 3.75 N of force. The blades that we chose do not consistently cut paper because of the blade profile. In future work, we intend to modify the blades to use a single-tapered scissor-like profile to improve cutting ability. However, all three sets of blades reliably and consistently cut through our target leaves. Therefore, we chose not to increase the strength of the motors, which adds weight, or to manufacture custom blades, which increases the system complexity.

D. Storing

The sample, once severed, is stored in a small box (Fig. 7a). The lid of the box closes, preventing non-target pathogens from contaminating the leaf. The box-shaped cavity can store clumps of leaves, small flowers, or pieces of stem.

E. Hardware and Software

We used lightweight and inexpensive servos for actuation. The lid is controlled by a Tower Pro SG90 servo with 1.8 kg-cm of stall torque. The blades are actuated with a pair of Futaba S3102 servos with 3.7 kg-cm of stall torque. The arm is controlled by a JR DS8411 servo with 11.2 kg-cm of stall torque.

To power the UAV, we use a 4-cell 4000mAh Multistar lithium polymer battery, and for the gripper a 3-cell, 1300 mAh Turnigy battery attached to a Drok 5 amp regulator set to 5 volts. The gripper battery and regulator weigh 132 g combined.

A DX7 transmitter was paired with a Spektrum AR8000 receiver to control the servos on the arm and gripper assembly. The wires for the servos have been shielded with aluminum foil to avoid interference between the gripper system and the flight system.

The UAV uses a Pixhawk flight controller running the open-source PX4 flight stack [19]. Flight status (battery voltage, flight mode) was monitored using QGroundControl graphical user interface.

III. Experimental evaluation

Two sets of experiments were carried out – with a standalone gripper separate from the flight system to quantify the range of grasping and cutting actions, and manual indoor flights to evaluate the success rate of leaf sample acquisition from a citrus tree. Citrus trees were chosen because of their relevance to early disease detection, but we attached other types of leaves from cutting trials on some flights to directly compare our two sets of experiments. In the rest of this section, we will describe the following results:

- Ability to grasp small and large leaves at varying angles.
- Ability to grasp small, leaf-like samples at varying depths and heights.
- Ability to cut through leaves and petioles.
- Performance of the system during indoor manual flight.

The standalone gripper tests were deemed successful if they indicated real-world feasibility of the system. Although this is hard to quantify for the non-uniform structures of leaves and plants, the experiments had to show success over a wide range of approach paths, angles, and contact areas. The manual flight tests were deemed completely successful if the system could repeatably approach, contact, grasp, sever, and securely store a sample before landing.

A. Grasping – leaf size and shape

In order to quantify how well the system could grasp leaves, we compared the angle of a leaf relative to the gripper to the number of successful grasping trials. At each angle, we adjusted the gripper height to ensure that the gripper could physically grasp the leaf before we began our 10 trials.

This test provided two key insights. First, as desired, our system has very steep cutoffs; if a leaf can be grasped once, it can be grasped reliably. Every set of 10 trials with small leaves had either 0 successes, or at least 7 successes. This experiment demonstrates that with a properly calculated approach trajectory, the gripper will succeed with high probability.

Second, the shape of the leaf does not matter significantly. We tested with two differently shaped, but similarly sized, leaves; one was long and pointed, the other was short and round (Fig. 13). The system grasped both types everywhere from 75 degrees below the horizontal to 75 degrees above the horizontal (Fig. 9). The two lower extents were limits of the experimental setup (Fig. 10). The round leaf was also more rigid than the pointed leaf, but this did not impact the performance of the gripper.

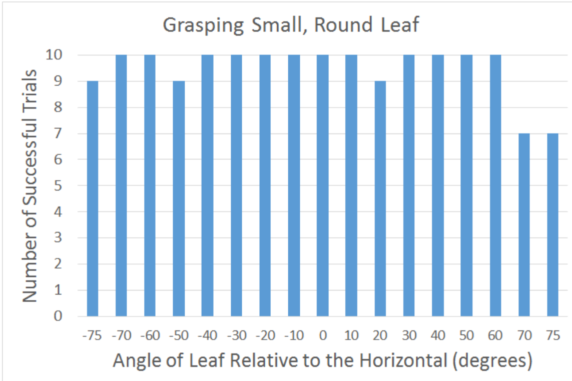

Fig. 9: Ability to grasp small, round leaf over 10 trials. The shape of this chart is representative of the other angle measurement tests performed.

Fig. 10: Setup of angle variation test experiment. Boxes were used to vary the height of the gripper so that the leaf could be grasped at the angle measured with the protractor.

Next, we investigated to see if large leaves could be grasped over the same angular range as small leaves. Our tests demonstrated that a large leaf could be grasped at angles ranging from -75 to 60 degrees relative to the horizontal. This range is less than the range for small leaves because of geometry; the stems of the large and the small leaves are held at the same angle, but the tip of the large leaf is further away from its stem, so the edge of that leaf sticks up past the reliable grasping zone. These tests demonstrate the robustness of our system across multiple leaf geometries.

B. Grasping – leaf angle, height, and depth

Leaves on trees will not have a specific angle of growth; they will show distributions of angle values. Our second experiment addresses feasibility of use in real-world settings by studying the relation between leaf angle, depth, and height in successful sample acquisition. To avoid the bias of a particular type of leaf and maintain precise measurements, we used a piece of printer paper marked in 0.5 cm increments for this test (Fig. 11). For each angle, we marked a new piece of paper and clamped it so that 4 cm were exposed. Measuring from the top of the storage mechanism, we varied the height by 0.5 cm increments from 1 cm until the paper could not be grasped. For each height, we placed the paper 1, 2, and 3 cm into the grasping mechanism. The angle varied from 0 to 50 degrees in 10 degree increments.

At each point, we recorded the largest covered increment. For example, if a paper was grasped so that the 1.5 cm mark was covered but the 2 cm mark was visible, we recorded 1.5 cm as our measurement. If the edge of the grasping mechanism was on a mark, we considered that mark covered. Unsurprisingly, our system's ability to grasp leaves decreased when less of the leaf was in range of the grasping mechanism, in both the depth and height dimensions. Our mock leaves were similar in size to the small leaves we tested with. The results for 20 degrees above the horizontal are shown in figure 12.

Fig. 11: Setup of depth test experiment.

Our experimental results demonstrate that the system can grasp leaves that enter the gripper mechanism at a large range of angles, heights, and depths, highlighting feasibility for real-world use. Due to environmental or other perturbations when approaching a leaf, the robot may not remain in line with the plant sample. However, our tests demonstrate that even small leaves can be grasped as long as 1–3 cm of the leaf enter the system. Our system held 1 cm of exposed leaf well for 0–10 degrees. It held 2–3 cm of exposed leaf well for 0–40 degrees. Large leaves can enter further into the gripper, allowing for a firmer grasp; this implies that the system will also work reliably on large leaves. We consider these tests successful because the system grasps leaves securely across a wide range of realistic conditions.

C. Cutting – leaves and petioles

Our third set of experiments tested the ability of our system to cut leaves. We tested four different types of leaves – two small and two large (Fig. 13). For both sizes, one type was flexible and thin, while the other was tougher. The large, tough leaves were also thicker, with a skeletal vein system. We tested all three pairs of blades individually on the small leaves. With the large leaves, all three blade pairs worked together to cut.

All three sets of blades consistently sever leaves completely. For the small leaves, every test resulted in either a complete sever or a partial sever. When the blades did not cut completely through the leaf, the remaining part could still be easily ripped off, indicating that the robot would successfully acquire and store the sample when it left the collection site.

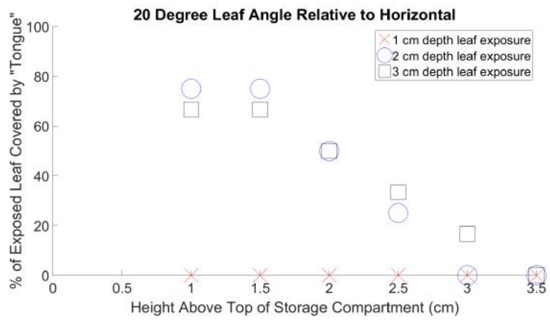

Fig. 12: Results of depth test with a paper leaf sample (for repeatability). The y-axis is normalized based on the depth of emulated leaf exposure. The linear trend shows that the ability to grasp leaf samples drops off as they move farther away from the storage box along the gripper's vertical axis.

Fig. 13: Leaves used in our tests. From left to right: large thick, large soft, small pointed, small round.

For large, soft leaves, our system severed the sample completely in almost every trial; again, in the trials with partial severing, the leaf easily tore free when the gripper was pulled away. Overall, for target leaves, the system succeeded in 87/90 trials. Only the set of tests using large, thick leaves showed decreased performance. In these trials, the blades did not cut completely through the leaf. For most of the trials, the leaf could still be easily removed even though the cut was not complete. The left, center, and right blades (looking down on the gripper) had a success rate on the thick leaves of 3, 5, and 7 out of 10 trials respectively. However, these leaves do not simulate target leaves, and the thickness of the leaf likely offset the precise spacing of the blades required for reliable cutting.

Our final manual cutting test dealt with a more purposeful application - rose cutting. First, we repeated our same cutting tests on rose leaves. As expected, all three sets of blades had a 100% success rate, with ten trials per blade set. The system was subsequently tested on rose petioles. Again, our mechanism had a 100% success rate over ten trials with each set of blades (60/60 total for leaves and petioles).

D. Flight Tests

We set up a lime (citrus) tree in our indoor lab space to complete these experiments. The tree was about 2 m tall. A pilot manually flew the quadrotor so that the leaf entered into the storage area of the gripper (Fig. 14). The system then cut the sample (Fig. 15), at which point we landed and repeated the test. The flights were conducted by the same pilot on two separate days.

After we collected data for citrus leaves, we attached small, pointed leaves and large, thin leaves (Fig. 13) to the tree using cable-ties. We wanted to ensure that our system could grasp multiple types of leaves from a realistic plant structure.

Fig. 16: Leaf cut by gripper. Section still on tree and section removed both shown.

We tested multiple UAV approach angles and quantified the sample acquisition success rate. In terms of multiple approach trajectories, we selected candidate leaves hanging at a variety of angles. They were located on branches spaced around the tree, so the pilot had to approach from different angles and heights while avoiding varying amounts of background foliage. The success rate was quantified in two ways. We examined the success rate of removing a specimen once the sample entered the gripper (Fig. 16); also, we measured the success rate of piloting a leaf into the gripper. We observed that UAV downwash did not affect leaf acquisition in any of our trials, demonstrating that we successfully designed for this interaction. Figure 18 summarizes flight results.

The results of the flight trials confirm our phytobiopsy UAV's usability. There are three specific potential failure points: failure to contact the leaf, failure to cut the leaf, and failure to store the leaf. With purely manual flight to a citrus tree, our pilot maneuvered a leaf into the grasping mechanism in 16/21 (76%) of the trials. Of those 16 trials, 15/16 (94%) instances resulted in the leaf being removed. One outlier occurred because the sample grasped was a mature, thick leaf and the blades were dull after excessive use. Finally, 12/15 (80%) of the successful cuts resulted in the sample securely sitting in the collection mechanism. In the case where the system did not properly cut the leaf (trial six), the sUAS crashed. However, once implemented, an autonomy stack will minimize the likelihood of similar errors. Fail-safes will be implemented to force the system to release and attempt a new approach trajectory.

At the end of first collection day (trial 11 in Fig. 18), we observed significant plant matter build-up on the blades, so we replaced them with new units. The cutting action immediately improved. Specifically, we tested the ability to cut the tough leaves that caused the crash and observed that the system could now sever those leaves (Fig. 17). On day two, we experienced no crashes and no instances of partial leaf severing. Both days had similar levels of storage success.

With the small, pointed leaves and the large, thin leaves (Fig. 13) attached to the tree, the system cut the sample successfully in 6/6 trials in which the leaf entered the system. Four of these trials were complete successes. One trial ended with the leaf trapped in the

Fig. 14: Images showing contact during 6 successful manual indoor flight trials on the citrus tree, in different configurations. Note that the effect of downwash is absent, due to the extended arm design.

Fig. 15: 7 leaf samples collected during the indoor trials shown in Fig. 14

blade mechanism instead of the collection mechanism. One trial ended with a successful collection, but the system could not be stabilized after the leaf suddenly cut (i.e. upon severing the leaf, the reaction forces changed suddenly, causing a crash). The system would have likely recovered had full autonomy been implemented. Additionally, this unusual cut occurred on day one with our dull blades. Even though our system did crash, it sustained minimal damage, demonstrating physical robustness. The sample did not enter the gripper in one additional small leaf trial. Overall, we demonstrate our phytobiopsy system's feasibility in a few ways. First, in 12/16 (75%) trials in which a sample entered the system, the system successfully removed and retained a leaf inside the collection mechanism. The pilot caused a large amount of variability in our manual tests, implying that performance will only increase with autonomy. Second, our tests had essentially the same success rate across two separate testing days (replacing the blades for day two actually improved the cutting reliability). Third, despite two crashes, our lightweight system sustained only minor damage to the arm. Fourth, our system successfully collected samples from multiple types of leaves that varied in thickness and size. These leaves hung at different, natural angles, heights, and locations relative to a citrus tree structure.

Of the above numbers, the 12/16 trials is the most important. This paper focuses on the design of the arm and gripper system for collection and storage of leaves. Although contacting the leaf is a necessary first step, this work deals more with the post-contact mechanism design. Therefore, the rate at which a specimen can be severed and stored is the system's critical measure of success.

IV. Conclusions and Future Work

This paper describes a novel, low-cost aerial phytobiopsy system, for robust and reliable leaf sample acquisition for ex-situ analysis. We discussed the mechanical design of the UAV, a mouth-like gripper assembly, extendable arm, and the hardware and software components. Extensive experimental validation quantified the gripper assembly's ability to grasp and store leaf samples. For small leaves, the system has a 150 degree range, and for large leaves it has a 135 degree range. We then demonstrate that with thin leaves and with rose petioles, the two target applications, our system has a 98% success rate for collecting samples by hand. Finally, we conducted a series of successful manual flight tests to show that not only is our system robust and reliable, but that it is feasible in real-world applications.

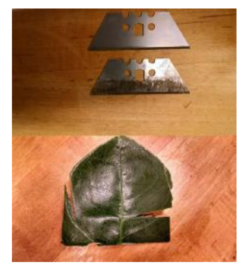

Fig. 17: Top: The new blade compared to the blade with plant matter build-up. Bottom: The partial cut was from an attempt with the old blades; the compete cut was from an attempt with the new blades.

Now that we have a functional prototype, the next step is autonomy, involving autonomous flight, detection, manipulation, and verification. The gripper assembly is already equipped with a 752 x 480 global-shutter RGB camera to provide visual feedback for the cutting action and for the approach path. Additionally, the platform must be able to visually determine symptomatic leaves, and the optimal approach trajectory in the presence of perturbations from wind as well as noise in state-estimates. A controller must be developed to compensate for the change of the center of mass when the arm extends. The system can be made lighter by making the arm out of a thinner, stiffer material like carbon fiber. Work on improving the blade profile is currently underway. Finally, we would like to see the system be able to collect multiple samples on one flight without cross-contamination.

To fully transition this system to an outdoor setting, we need to verify that it can detect and interact with leaves while rejecting disturbances such as differing light levels, wind, and occluding leaves. Ultimately, we envision this system being incorporated in heterogeneous robot teams (UAVs or ground robots depending on farm setting) that can autonomously search for crop diseases using vision algorithms and machine learning, and then collect samples of the diseased plant and transport them to a nearby lab for analysis.

	Day One Trials										
Trial Number	1	2	3	4	5	6	7	8	9	10	11
Contacted Sample		X	X	X		X	X	X	X	X	X
Removed Sample		X	X	X			X	X	X	X	X
Retained Sample		X	X	X			X		X		X

	Day Two Trials										
Trial Number	12	13	14	15	16	17	18	19	20	21	Summary
Contacted Sample			X	X	X	X	X	X		X	16/21
Removed Sample			X	X	X	X	X	X		X	15/21
Retained Sample			X	X		X	X	X		X	12/21

Fig. 18: Manual Flight Test Results: 1-8 and 12-17 were citrus leaves. 9-11 and 19-21 were small, pointed leaves attached to citrus tree, and 18 was a large, thin leaf attached to the same tree. Trial 6 resulted in a crash from a toughened leaf, with no system damage. Trial 11 resulted in a crash from changes in reaction forces after collection, causing minor arm damage. Trial 15 saw significant tearing at the sample interface. The specimen became stuck in the blades instead of the collection chamber in trials 8, 10, and 16.

V. Acknowledgements

We gratefully acknowledge NSF grant CNS-1521617 and USDA grant 2015-67021-23857 under the National Robotics Initiative for supporting this work. We thank Samara Gray and Josh Darfler of the University of Pennsylvania greenhouse for keeping the citrus tree in good health.

References

[1] S. Sarkar, J. Das, R. Ehsani, and V. Kumar, "Towards autonomous phytopathology: Outcomes and challenges of citrus greening disease detection through close-range remote sensing," in *IEEE International Conference on Robotics and Automation*. IEEE, 2016, pp. 5143–5148.

[2] K. Ahlin, B. Joffe, A.-P. Hu, G. McMurray, and N. Sadegh, "Autonomous Leaf Picking Using Deep Learning and Visual-Servoing," in *5th IFAC Conference on Sensing, Control and Automation Technologies for Agriculture (AGRICONTROL)*, vol. 49, no. 16, 2016, pp. 177 – 183.

[3] D. Jenkins and B. Vasigh, "The Economic Impact of Unmanned Aircraft Systems Integration in the United States," *Association for Unmanned Vehicle Systems International Economic Report*, vol. March, 2013.

[4] J. Das, G. Cross, C. Qu, A. Makineni, P. Tokekar, Y. Mulgaonkar, and V. Kumar, "Devices, systems, and methods for automated monitoring enabling precision agriculture," in *2015 IEEE International Conference on Automation Science and Engineering*, Aug 2015, pp. 462–469.

[5] J. Thomas, J. Polin, K. Sreenath, and V. Kumar, "Avian-Inspired Grasping for Quadrotor Micro UAVs," in *ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE)*, 2013.

[6] M. Luo, W. Lu, B. Sun, and T. Mei, "Autonomous Grasping of a Space Robot Multisensory Gripper," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2006, pp. 1014–1019.

[7] J. Thomas, G. Loianno, K. Daniilidis, and V. Kumar, "Visual servoing of quadrotors for perching by hanging from cylindrical objects," *IEEE Robotics and Automation Letters*, vol. PP, no. 99, pp. 1–1, 2015.

[8] G. Garimella and M. Kobilarov, "Towards model-predictive control for aerial pick-and-place," in *2015 IEEE International Conference on Robotics and Automation*, May 2015, pp. 4692–4697.

[9] S. Kim, S. Choi, and H. J. Kim, "Aerial manipulation using a quadrotor with a two DOF robotic arm," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov 2013, pp. 4990–4995.

[10] K. Alexis, G. Darivianakis, M. Burri, and R. Siegwart, "Aerial robotic contact-based inspection: planning and control," *Autonomous Robots*, vol. 40, no. 4, pp. 631–655, 2016.

[11] F. Huber, K. Kondak, K. Krieger, D. Sommer, M. Schwarzbach, M. Laiacker, I. Kossyk, S. Parusel, S. Haddadin, and A. Albu-Schäffer, "First analysis and experiments in aerial manipulation using fully actuated redundant robot arm." in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 3452–3457.

[12] C. Papachristos, K. Alexis, and A. Tzes, "Technical Activities Execution with a TiltRotor UAS employing Explicit Model Predictive Control," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 11 036 – 11 042, 2014.

[13] P. E. I. Pounds and A. M. Dollar, " Stability of Helicopters in Compliant Contact Under PD-PID Control," *IEEE Transactions on Robotics*, vol. 30, no. 6, pp. 1472–1486, 2014.

[14] J.-P. Ore, S. Elbaum, A. Burgin, and C. Detweiler, "Autonomous Aerial Water Sampling," *Journal of Field Robotics*, vol. 32, no. 8, pp. 1095–1113, 2015.

[15] D. V. Gealy, S. McKinley, M. Guo, L. Miller, S. Vougioukas, J. Viers, S. Carpin, and K. Goldberg, "DATE: A handheld co-robotic device for automated tuning of emitters to enable precision irrigation," in *2016 IEEE International Conference on Automation Science and Engineering*, Aug 2016, pp. 922–927.

[16] K. Sreenath, T. Lee, and V. Kumar, "Geometric control and differential flatness of a quadrotor UAV with a cable-suspended load," in *52nd IEEE Conference on Decision and Control*, Dec 2013, pp. 2269–2274.

[17] F. Ruggiero, M. A. Trujillo, R. Cano, H. Ascorbe, A. Viguria, C. Peréz, V. Lippiello, A. Ollero, and B. Siciliano, "A multi-layer control for multirotor uavs equipped with a servo robot arm," in *2015 IEEE International Conference on Robotics and Automation*, May 2015, pp. 4014–4020.

[18] M. Mahvash, L. M. Voo, D. Kim, K. Jeung, J. Wainer, and A. M. Okamura, "Modeling the Forces of Cutting With Scissors," in *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 3, 2008, pp. 848–56.

[19] L. Meier, D. Honegger, and M. Pollefeys, "PX4: A node-based multithreaded open source robotics framework for deeply embedded platforms," in *2015 IEEE International Conference on Robotics and Automation*, May 2015, pp. 6235–6240.